

# Leveraging Crowd for Game-based Learning: A Case Study of Privacy Education Game Design and Evaluation by Crowdsourcing

Wendy Wang<sup>1</sup>, Yu Tao<sup>2</sup>, Kai Wang<sup>3</sup>, Dominik Jedruszczak<sup>1</sup>, Ben Knutson<sup>1</sup>  
 Department of Computer Science<sup>1</sup>, the College of Arts and Letters<sup>2</sup>, School of Management and  
 Marketing<sup>3</sup>  
 Stevens Institute of Technology<sup>1,2</sup>, Kean University<sup>3</sup>  
 Hoboken, NJ<sup>1,2</sup>, Union, NJ<sup>3</sup>  
 hwang4,ytao,djedrusz,bknutson@stevens.edu<sup>1,2</sup>, wangkaicv@gmail.com<sup>3</sup>

## ABSTRACT

As the Internet grows in importance, it is vital to develop methods and techniques for educating end-users to improve their awareness of online privacy. Web-based education tools have been proven effective in many domains and have been increasingly adopted by many online professional and educational services. However, the design and development of Web-based education tools for online privacy is still in the early stage. The traditional solutions always involve privacy experts who have sophisticated expertise. Such involvement can make the tool development costly. Furthermore, it is not clear how inspiring and effective these education tools are to general users of varying backgrounds, specially to novice users who have rarely dealt with online privacy issues before. In this paper, we design, develop, and evaluate a game-based privacy learning system by leveraging the wisdom of a crowd of non-experts on Amazon Mechanical Turk. Empirical study demonstrates that the crowd can provide high-quality ideas of designing and developing a practical, educational privacy learning game.

## 1. INTRODUCTION

Online social networking communities have undergone an explosion in recent years, as both the kinds and numbers of sites have grown and their memberships increased. This has led to the proliferation of personal data revealed by users of online communities, which presents a variety of privacy risks. Internet users still know little about online privacy [4], even though their awareness of online privacy has increased [13]. Due to the fact that the existing privacy education tools are either too technical or impractical [24], there is a need for new, effective privacy education tools for Internet users of all ages and backgrounds.

In this paper, we consider the privacy education that takes the format of game-based learning. Game-based learning is the use of digital games with educational objectives. It pro-

vides effective, interactive experiences that motivate users to actively engage in the learning process when playing with a game [28]. However, developing an effective privacy learning game is challenging. The traditional solutions always involve privacy experts who have sophisticated expertise, deep technical skills, and broad experience. However, the involvement of professionals can make the game development costly. Furthermore, it is not clear how motivating, inspiring, and/or effective these education games are to general users of varying backgrounds, experience, and expertise of privacy, especially to novice users who have rarely dealt with privacy issues before.

Besides practicality, creativity is another important design factor of an effective game-based learning system [33]. Facilitated by advances in web-related technologies, crowdsourcing has become a new technique for practicing open innovation, for example, the sheep market [2], T-shirt design [9], and software idea competition [21]. It has shown that crowdsourcing workers indeed can perform some creative tasks to reduce the production costs [22, 31].

Besides the development of creative systems, crowdsourcing also has been proven effective for the evaluation of system performance. Examples of evaluation tasks include relevance evaluation for information retrieval [5], book search evaluation [15], and inference rule evaluation [37]. Few work has used crowdsourcing to evaluate the effectiveness of a game-based learning system.

The aim of this paper is to explore the potential of using a crowd of non-experts to help design and evaluate a privacy-education game whose target players are Internet users of all ages and backgrounds. There are several challenges of utilizing crowdsourcing for both design and evaluation of a game-based privacy learning system. First, a crowdsourcing system (e.g., Amazon Mechanical Turk [1]) typically supports microtasks that are simple, fast-paced, and require the least amount of training. It is not clear whether the crowd can perform the task such as creating a privacy education game that is complex, time-consuming, and requires some background knowledge in online privacy. Second, the crowd workers have various backgrounds and knowledge of online privacy. They may create game design ideas that are dramatically different in terms of both content and format. It is difficult to aggregate these ideas with significant individual heterogeneity. Third, it is challenging to evaluate the quality of the crowd's ideas, as well as the effectiveness and educational impact of a privacy education game, with quan-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Copyright 20XX ACM X-XXXXX-XX-X/XX/XX ...\$15.00.

titative measurements.

To address these challenges, we develop a prototype of *DEPET*, Development and Evaluation of Privacy Education Game. The system consists of three phases, *idea collection*, *game development*, and *game evaluation*. In the *idea collection* phase, *DEPET* collects and aggregates the ideas for designing a privacy education game from a crowd of non-experts on Amazon’s Mechanic Turk. In the *game development* phase, the privacy education game is implemented based on the aggregated ideas from the *idea collection* phase. In the *game evaluation* phase, the developed privacy game is released to both the crowd and privacy experts for effectiveness evaluation. The evaluation feedback information is then used to further improve the game.

Our contributions include the following. First, we are the first who exploit crowd innovation for the game-based privacy learning. Second, we investigate the main characteristics of a wise group that is capable of developing creative and high-quality ideas for the development of an effective privacy education game. Third, we develop a prototype of privacy education game that instantiates the crowd’s ideas. Fourth, we develop various metrics for assessing the quality of the crowd-generated ideas in terms of correctness and novelty. Fifth, we leverage both privacy experts and a large crowd of non-experts on Amazon’s Mechanic Turk to evaluate the quality of the privacy education game with regard to its usability and educational impact. The evaluation results show that crowd can provide high-quality inputs for both creative system design and system evaluation.

The paper is organized as following. Section 2 defines the scope of tasks that are to be assigned to the crowd. Sections 3, 4 and 5 explain the details of the *idea collection*, *game development* and *game evaluation* phases respectively. Section 6 presents the details of experiments. Section 7 discusses the related work. Section 8 concludes the paper.

## 2. DEFINITION OF PROBLEM SCOPE

In this section, we define the scope of the privacy education game. The scope is defined at two dimensions: (1) the *privacy* dimension which addresses the privacy issue that the education game targets at, and (2) the *design space* dimension which defines the game design elements desired from the crowd. Next, we discuss the details of these two dimensions.

### 2.1 Privacy Domain

Internet privacy is a broad term that refers to a variety of factors, techniques and technologies used to protect sensitive and private data, communications, and preferences. An education game that covers such a broad concept is expected to be complicated and cumbersome. Therefore, in this project, we narrow down to a specific privacy domain. Which privacy domain to pick is decided by two factors: (1) whether the picked domain can lead to an effective privacy education game to general users; and (2) whether the picked domain is suitable for the crowdsourcing task that requires open innovation design ideas. Taking both factors into consideration, we pick a specific privacy issue that has received considerable concerns by Internet users: *privacy issues of social networking sites*. We pick this issue due to the following two reasons. First, a key requirement of a good game-based learning tool is to create an effective game-based learning environment that users are familiar with [28]. According to a recent survey, 74% of online adults use social networking

sites [3]. Therefore, a game that focuses on privacy issues of social networking sites meets the requirement. It should enable general users to enjoy the game by leveraging their real-life social networking experience. This also should enable general users to get highly engaged in the learning process and can easily transfer what they learned from the simulated game environment to their real life. This addresses Factor 1. Second, in general, designing game-based learning tools requires the designers to be equipped with knowledge or related experience. Common crowdsourcing workers are non-experts in privacy. However, given the popularity of social networks, we expect that non-expert crowd workers can contribute high-quality design ideas of the privacy learning game by leveraging their real-life experiences of using social networking sites. This addresses Factor 2.

### 2.2 Design Space

Game design is a complex process. It involves many components, including rule development, design of story plot, grading and awarding strategy, just to name a few. A *design space* identifies the range of possible solutions for a design problem [10]. The definition of design space in the crowdsourcing task is important for the quality of the collected results. A design space that is too broad and open-ended may not be able to harvest high-quality ideas, while the tasks that are too closed-ended may leave too little design space for the workers to generate innovative ideas. Due to the fact that the types of tasks accomplished through typical crowdsourcing systems (e.g., Amazon Mechanical Turk) have been limited to those that are low complexity, independent, and require little time and training to complete [18], asking the workers to design a whole privacy education game from the scratch is too difficult to accomplish. Therefore, we define the game rules and GUI interface by ourselves, and ask the workers to fill in the game content only. More details of the game content will be discussed in Section 3.

## 3. IDEA COLLECTION AND ANALYSIS

In the *Collection* phase, we generated Human Intelligence Task (HIT) on Amazon Mechanic Turk (AMT), asking workers to contribute ideas of the content of a privacy education game. In particular, the required game content includes:

1. Examples of postings shared on social networking sites (e.g., Facebook) that may expose private and sensitive aspects of one’s personal life; and
2. Suggestions of how to fix the postings to remove the privacy information.

In the HIT, each worker read a preview consisting of the instructions and an example of privacy game interface. The preview was the same for all workers. The workers’ answers were evaluated by the privacy experts. All workers received the payment of \$0.10 after they submitted their ideas.

### 3.1 HIT Design

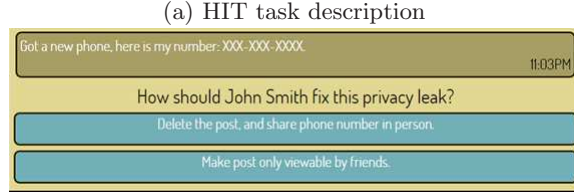
Our HIT task description consists of three parts:

- In *Preview* part, we describe the purpose of the task, and the rewarding strategy.
- In *Instruction* part, we present the detailed instructions of the task.

Please answer the following questions.

What personal information on Facebook needs to be protected?

Please provide three scenarios where this personal information is not well protected. Also please propose some solutions to fix the situation and ensure proper protection of the information. Please be creative and provide unique answers.



(b) An example of online posting used in HIT

Figure 1: HIT of the open-ended experiments

- In *Background* part, we collect the workers’ demographic information and their background knowledge in online privacy.

We designed two HIT tasks under different settings, namely the *open-ended* setting and the *closed-ended* setting.

**Open-ended setting.** By this task, we asked the workers to contribute any example of personal information on Facebook that needs to be protected. Each worker is asked to provide three examples of online postings that revealed some sensitive personal information. Each workers is also asked to provide fix solutions for these examples, one solution for each example. In HIT description, we emphasize *creativity* as one of the requirements of the examples. Figure 1 (a) shows a part of instructions that we used for the open-ended setting. To help the workers understand the task requirement and motivate them to contribute high-quality ideas, in the instructions, we include an example of online posting that reveals a person’s private phone number. We also include two possible fix solutions of the problematic posting in the example. Figure 1 (b) shows the example that we used in the HIT instruction.

**Closed-ended setting.** By this task, we focus on *four specific* privacy topics: (1) medical information, (2) income, (3) work history, and (4) student records, and asked the workers to contribute a number examples of online postings that reveal privacy information of any of these four types. We picked these four specific types of privacy information as these are the ones that the results of the *open-setting* experiments did not cover at all (more analysis of the *open-ended* experiment results can be found in Section 3.2).

The task descriptions of the open- and closed-ended settings have the same *Preview* and *Background* parts; they only differ at the *Instruction* part. Due to the disjoint types of privacy information that were covered by both open- and closed-ended experiments, we allow the same workers to participate in both experiments. Both open- and closed-ended settings have the same rewarding mechanism.

**Workers’ background.** As a part of HIT description, we collected the workers’ demographic information, including their locations, gender, and education level. We also surveyed their background knowledge in online privacy, includ-

ing: (1) how often do they use Facebook; (2) whether they have adequately protected their own private information on Facebook; and (3) how much do they know about online privacy. We collect these information trying to find out the correlation between the quality of ideas and the workers’ background.

## 3.2 Idea Aggregation and Analysis

### 3.2.1 Idea Aggregation

Sixteen workers participated in the idea generation phase, each contributing three example postings that revealed some sensitive personal information. We observe that these forty-eight examples share a common set of privacy topics. Therefore, we design an idea aggregation method that enables to automatically aggregate the collected crowd-generated design ideas. A possible solution is to use topic modeling [25]. However, there are several challenges of using topic modeling to aggregate crowd’s ideas. First, many existing topic modeling methods (e.g., LDA [7]) rely on a large document corpus to get decent results, while the idea description collected from AMT are short (no more than 200 characters in our case) and abbreviated. The number of ideas is also small (48 examples), given our limited budget to hire AMT workers. Second, the crowd’s input is noisy - it contains typing errors and meaningless words. Due to these two challenges, applying the existing topic modeling methods to the crowd’s ideas fails to deliver good results. Therefore, instead of using any topic modeling method based on training and learning, we use the keyword extraction method based on word frequency analysis. We use TF-IDF to weight the terms. Then we identify those terms that are appropriate as privacy topics. We found thirteen distinct privacy topics by this method. We grouped the examples provided by the crowd according to the privacy topics that they are associated with. It turned out that the privacy topics *activity*, *home address*, and *personal emotions* are the three that are associated with the highest number of examples, while the popular privacy topics such as *name*, *age*, and *birth date* has fewer examples. This is surprising as the privacy concerns of revealing activities and personal emotions on social networking sites are indeed paid much less attention as name and birth date [14]. This shows that the workers indeed intended to contribute creative ideas of privacy examples, as required in the HIT instructions.


### 3.2.2 Idea Analysis

In this subsection, we analyze the aggregated ideas for both open-ended and closed-ended settings.


**Open-ended Setting.** Sixteen workers participated in this round of experiments, each contributing three example postings that revealed some sensitive personal information. There are forty-eight examples in total. We analyze these examples and have the following interesting observations.

First, there exists a large overlap on the type of *privacy topics* that the examples covered. We categorized the examples by their privacy topics, and counted the number of examples that each privacy topic has. There are thirteen distinct privacy topics in total. Figure 2 (a) shows these privacy topics and the number of examples that each privacy topic covers. Among these topics, *activity*, *home address*, and *personal emotions* are the three topics that are of the highest frequency, while the popular privacy topics such as name, age, and birth date has fewer examples. This is surprising

as the privacy concerns of revealing activities and personal emotions on social networking sites are indeed paid much less attention as name and birth date [14]. This shows that the workers indeed intended to contribute creative ideas of privacy examples, as required in the HIT instructions.



(a) Number of examples of each privacy topic



(b) Number of examples for each population group

**Figure 2: Analysis of Crowd-generated Ideas**

Second, the crowd-generated example postings are very diverse in terms of the population groups that the examples can apply to. The examples that belong to the same privacy topic group indeed target on different group populations. For instance, in group of the privacy topic *location*, one example described an adult’s posting showing his/her exercise routine at a specific gym, while another example mentioned the location of a kid’s favorite playground. We categorize the workers’ examples into four population groups: (1) kids and high school students (age < 18); (2) college students (age in [18, 24]); (3) young professionals (age in [25, 30]); and (4) older adults (age > 30). We assign the examples to these four population groups, and count the number of examples in each group. The grouping of examples is not exclusive; some examples can be applied to multiple groups. Figure 2 (b) shows the grouping results. It turned out that the *older adult* group received the highest number of privacy topics. This may be due to the fact that a majority of workers are in the same age range as the older adult group; they design the examples from their real-life experiences.

**Closed-ended Setting.** In this round, twelve workers participated, each creating three example postings. There are thirty-six examples in total. These examples cover four specific types of online privacy topics (as required in HIT), namely medical information, income, work history, and student records. Each topic has nine examples. We found out that none of the examples are similar. More details of the

similarity measurement and the analysis of the example similarity can be found in Section 6.4. Therefore, we aggregate the ideas collected from the closed-ended setting based on the four privacy topics.

## 4. GAME DEVELOPMENT

In this phase, we implemented a prototype of our Web-based privacy education game *DEPET*<sup>1</sup>. The ideas that were collected and integrated from the crowd were implemented as Facebook-style posting examples in the game. *DEPET* was implemented by using MySQL, Javascript, and HTML5.

According to the four different target groups of population that the crowd-generated example postings covered, we designed four different characters in the game. Each character is associated with a set of privacy topics. The privacy topics that each character is associated with are:

- Sophia (high school student, age < 18): ID, credit cards, driver license, and academic behaviors;
- Aiden (college student, age [18, 24]): home address, academic behaviors, medical information, and personal emotions;
- Olivia (young adult, age [25, 30]): professional issues;
- Lucas (older adult, age > 30): information of family members, personal ideology, and medical information.

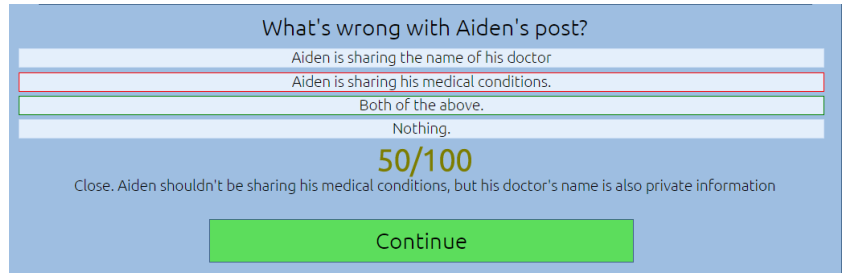
We design the following character-based game rules. The players can pick any of the four aforementioned characters to start the game. Each game character has five or six online postings, each posting revealing some private and sensitive aspects of one’s personal life. Each posting is also associated with three possible fix solutions. Most of the example postings and the fix solutions in the game are generated from the *idean collection* phase. The players have to: (1) decide whether these postings have privacy issues, and what these privacy issues are if there is any; and (2) what is the best fix solution in terms of balanced privacy and data availability. Users’ different decision of where the privacy problem occurs is assigned with a different performance score. An example of the postings is shown in Figure 3 (a). Besides showing the scores, the game gives the players the feedback of their choices. The feedback advises the players why their choices are (in)correct, and what are the real privacy problems of the postings. An example of the feedback is shown in Figure 3 (b). The feedback was designed by the privacy experts. After the players finishing judging the privacy problems of the postings, the game shows three to four fix solutions. An example of the fix solutions is shown in Figure 3 (c). The best solution is the one that best addresses the trade-off between privacy and data availability. The players pick one solution. Each solution has a score, depending on how good privacy and data availability is balanced. For instance, consider the four solutions shown in Figure 3 (c). Removing the image alone (Option A) or only the specific medical condition in the text (Option B) cannot protect the privacy entirely. Therefore, the only correct solution is to remove the post. The game also gives the feedback to the players of their picked fix solutions. The feedback states clearly why the picked solution can(not) remove the privacy leak sufficiently.

<sup>1</sup>The prototype is available at <http://stevens.edu/depot>.





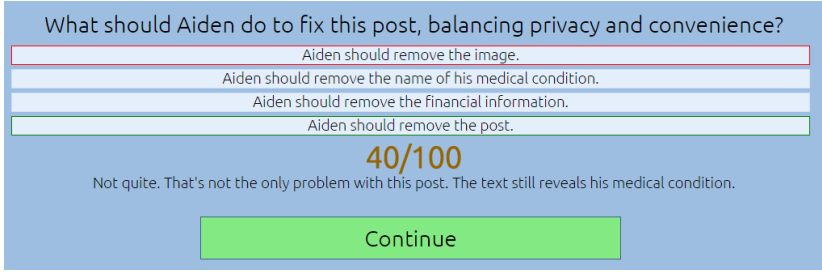
(a) An example of posting that has privacy issues



(b) An example of the feedback of privacy issues



(c) An example of fix solutions



(d) An example of the feedback of picked solution

Figure 3: Example Posting and Possible Fix in Game

## 5. EVALUATION BY CROWD

In this project, we evaluated the *DEPET* prototype in terms of its usability and education success by leveraging the crowd on Amazon Mechanical Turk.

### 5.1 Task Description

Our task description consists of three parts:

- In *Preview* part, we described the purpose of the evaluation task and the rewarding strategy. Each participant is paid \$0.20 after they finish the evaluation.
- In *Instruction* part, we present the detailed instructions of the task. The Web link of *DEPET* is attached in the instructions.
- In *Background* part, we collected the evaluators' demographic information, including their locations, gender, and education information. We also surveyed their background knowledge in online privacy, including how often do they use Facebook and their self-evaluation of privacy levels.

To ensure that all four game characters are evaluated by the same number of workers, we released four evaluation HITS, one for each character. In each HIT, we asked the player to evaluate the character as required in the task description. We require that each participant reviews and plays with all example postings and the suggested fix solutions (of one specific character) before (s)he evaluates the game. To ensure this, we created a confirmation code for each character. Only those workers who finished all example postings of the character will receive the confirmation code. Different characters have different confirmation codes. The workers are asked to enter the confirmation code before they start the evaluation. Figure 4 shows the part of instructions with the confirmation code field.

The evaluation took the format of user surveys. After the workers enter the confirmation code, they answer a set of

#### Instruction page

Thank you for participating in this survey for an online privacy educational game. This game is still in an early version and we would like to have your evaluation and comments based on your playing of the game. Your input will make the educational game more efficient, effective, and user-friendly.

First, please play the game in the link below. Although there are four different characters in the game, please play with the character "Aiden". You will need to get a confirmation code after you finish with this specific character. You will be paid only after we verify the confirmation code.

#### Online Game

The questions below ask for your evaluation on a few aspects of a successful educational game, including learnability, performance effectiveness, flexibility, and user satisfaction. On a scale of 1 to 5, with 1 being strongly disagree, 2 being disagree, 3 being neutral, 4 being agree, and 5 being strongly agree, please select the number that best matches your experience.

Please provide the confirmation code.

Figure 4: Evaluation HIT task description

questions in the survey regarding the game character that they played with. The survey evaluates two types of effectiveness of the developed game prototype for privacy education and learning:

- *Tool usability*, which measures whether the tool (i.e., the game) is easy to use; and
- *Educational success*, which measures the impacts of the tool on delivery and learning of online privacy knowledge.

There are twenty questions in the survey. Among these questions, fifteen questions are single-choice questions, and five questions are free-text format. The single-choice questions ask the users to select the number that best matches their experience. The numbers specifies a scale of 1 to 5, with 1 being strongly disagree, 2 being disagree, 3 being neutral, 4 being agree, and 5 being strongly agree. The free text-

format questions ask the workers to enter their feedback and suggestions. Next, we discuss how we designed the questions in the survey.

**Evaluation of tool usability.** We used the *usability testing* method to evaluate tool usability. Usability tests, with a user-centered approach, have been successfully used in the development of other products such as interactive multimedia software and web-based learning tools [20, 32, 26]. The usability test model by [20] suggests to test: (1) Learnability (the ease of learning to use the system, e.g., clearly labeled components, easy navigation); (2) Performance effectiveness (the ease of using the system in terms of speed and error rate); (3) Flexibility (the level of freedom to use multiple commands to achieve the same goal); (4) Error tolerance and system integrity (the ability to recover from errors and prevent data corruption or loss); and (5) User satisfaction (users’ opinions and perception of the training tool). For each testing component, we designed corresponding questions in the HIT, asking for the crowd’s feedback.

**Evaluation of educational success.** Kirkpatrick proposed a training evaluation model [16] that can objectively analyze the effectiveness and impact of education. We apply Kirkpatrick’s evaluation model to our evaluation phase. Kirkpatrick’s evaluation model measures the following four items.

**Level-1 Reaction Measurement.** In the survey, we ask the questions related to what the workers thought and felt about the game. For example, did they feel that the game was worth their time? What were the biggest strengths and weaknesses of the game? Did they like the topic, the materials, and the format? By measuring reaction, we will understand how well the game-based learning was received by the workers.

**Level-2 Learning Measurement.** To measure how much the workers’ knowledge of online privacy has increased as a result of the education game, we asked the workers in the survey whether their knowledge of online privacy has increased as a result of the playing the game.

**Level-3 Behavior Measurement.** We ask a series of questions regarding how the workers would apply the knowledge learned from the game. For example, will the workers put any of their learning of online privacy techniques to their practical use? Are they able to teach their new knowledge and skills to other people? The main focus is the behavioral changes due to the game playing.

**Level-4 Results Measurement.** We analyze the effect of the game on the improved performance of the workers. Kirkpatrick suggested several key performance indicators, such as volumes, percentages, and other quantifiable aspects of performance, such as whether the number of privacy accidents has been reduced [16]. It is hard to collect such observable performance indicators in the survey. Therefore we did not ask such questions. Instead we asked the workers whether they would recommend this game to others for the learning on online privacy.

There are twenty questions in the survey. We categorize all the questions into three components:

- Evaluation of game interface, including seven questions that cover the evaluation of tool usability;
- Evaluation of game content, including seven questions that cover the evaluation of both tool usability and education success; and

- Overall evaluation, including six questions that cover the evaluation of education success.

## 6. EXPERIMENTS

In this section, we explain the details of our experiments.

### 6.1 Setting

We performed both the idea collection and game evaluation tasks on Amazon Mechanical Turk. Besides the crowd, we also have three privacy experts: a professor whose expertise in data privacy and two of her senior PhD students.

### 6.2 Workers’ Background

We collected the demographic information of workers, as well as their knowledge of online privacy, for both Collection and Evaluation phases.

#### 6.2.1 Idea Collection Phase

**Open-ended setting.** Sixteen workers participated this experiment. All users are from US. The mean of their ages is 39. Regarding the education level, ten workers did not finish high school. Four finished high school, and two had college education. Regarding the gender, ten workers are females and six were males. Regarding the background knowledge and experiences on social networking privacy, all workers use Facebook at least every month (Figure 5 (a)). Most of them considered themselves having protected their privacy on Facebook sufficiently (Figure 5 (b)). A majority of these workers considered themselves have intermediate level of privacy knowledge (Figure 5 (c)).

**Closed-ended setting.** Twelve workers participated this experiment. All users from US. The mean of their ages is 26. Regarding the education level, nine workers finished high school, and three received college education. Regarding the gender, seven workers are females and five were males. Regarding the background knowledge and experiences on social networking privacy, all workers except one use Facebook at least every month (Figure 6 (a)). Most of them considered themselves having protected their privacy on Facebook adequately (Figure 6 (b)). All these workers considered themselves have intermediate level (scale 3/4 out of 5) of privacy knowledge (Figure 6 (c)).

To summarize, in both settings, the crowd who generated game design ideas are of various demographic backgrounds. They also have different knowledge and experiences in online privacy.

#### 6.2.2 Game Evaluation Phase

One hundred and twenty workers participated into the evaluation task. All of them are from US. The average age is 33. Regarding the gender, sixty-two workers are females, and fifty-eight are males. Regarding the education level, 34 workers did not finish high school, 74 workers finished high school, and 12 workers received college education. Regarding their online privacy experience, 11 workers considered themselves did not protect privacy sufficiently, 85 workers considered them with sufficient protection of online privacy, and 24 workers were not sure. Regarding their privacy expertise, one worker knows nothing about online privacy, 5, 45, and 69 workers labeled themselves as level 2, 3, and 4 (out of 5) respectively. None of the workers considered themselves as an expert on online privacy. The average of self-evaluation expertise level is 3.69 (out of 5). Again, the




Figure 5: Privacy background and experiences of the workers in open-ended setting




Figure 6: Privacy background and experiences of the workers in close-ended setting

crowd who evaluated the game came from various demographic backgrounds, with different knowledge and experiences in online privacy.

### 6.3 Metrics of Idea Quality Evaluation

We define two metrics, namely the *correctness* and *novelty*, to evaluate the quality of the crowd-generated ideas.

**Correctness.** Informally, the correctness measurement of workers’ ideas is to verify whether the ideas indeed expose private and sensitive aspects of one’s personal life (explicitly or implicitly). We refer to the privacy experts’ evaluation for correctness measurement. Each idea receives a binary value (0/1) indicating its correctness. Given  $m$  ideas, the correctness ratio is defined as

$$cr = \frac{k}{m},$$

where  $k$  is the number of correct ideas. Intuitively, the correctness ratio measures the portion of the ideas that are correct.

**Novelty.** The novelty measurement is based on the *similarity* of ideas. To measure the pairwise similarity of ideas, first, we manually extract the keywords of each idea. We only kept the meaningful keywords with stems removed. Synonyms are merged to one representative word (e.g., “photo” and “picture” are represented by “photo”). Then for every two ideas  $E_1$  and  $E_2$ , let  $K_1$  and  $K_2$  be their keywords, we measure the similarity of  $K_1$  and  $K_2$  by measuring the similarity of  $E_1$  and  $E_2$ . There are a number of approximate string similarity measurement methods in the literature, e.g.,  $q$ -grams, edit distance, and Euclidean distance (See [19] for a good tutorial.) In this paper, we consider the *Jaccard* similarity based on  $q$ -grams. In particular, each keyword is decomposed into a set of tokens of length  $q$  (called  $q$ -grams). Then for any two strings  $S_1$  and  $S_2$ ,

$$jaccard(S_1, S_2) = \frac{|G(S_1, q) \cap G(S_2, q)|}{|G(S_1, q) \cup G(S_2, q)|}, \quad (1)$$

where  $G(S_1, q)$  ( $G(S_2, q)$ , resp.) is the set of  $q$ -grams of  $S_1$  ( $S_2$ , resp.). We say two keywords  $S_1$  and  $S_2$  are  $\delta$ -similar regarding the Jaccard metrics, if  $jaccard(S_1, S_2) \geq \delta$ . Given two sets of keywords  $K_1 = \{S_1^1, \dots, S_m^1\}$  and  $K_2 = \{S_1^2, \dots, S_n^2\}$ , let  $\ell$  be the number of  $\delta$ -similar keyword pairs of  $K_1$  and  $K_2$ , the similarity of  $K_1$  and  $K_2$  is measured as

$$sim(K_1, K_2) = \frac{\ell}{|K_1 \cup K_2|}. \quad (2)$$

We say two ideas  $D_1$  and  $D_2$  are similar, denoted as  $D_1 \approx D_2$ , if the similarity of their keywords  $K_1$  and  $K_2$  satisfies  $sim(K_1, K_2) \geq \theta$ , where  $\theta \in [0, 1]$  is a user-specified threshold.

Given a set of ideas  $D_1, \dots, D_n$ , the novelty of  $D_i$  ( $1 \leq i \leq n$ ) is measured as:

$$novelty(D_i) = \frac{n - |\{D_j | D_j \approx D_i, i \neq j\}|}{n}.$$

That is, the novelty of  $D_i$  measures the percentage of ideas that are dissimilar to  $D_i$ . Intuitively, the more the dissimilar ideas there are, the higher the novelty of  $D_i$  is.

### 6.4 Idea Quality

In this set of experiments, we measure: (1) both the correctness and novelty of the ideas collected by the Collection phase; and (2) the correlation between the quality of workers’ ideas and their privacy background knowledge and experiences.

**Correctness & Novelty.** First, we measure the correctness of the collected forty-eight posting examples and their suggested fix solutions in the open-ended setting. Seven examples were evaluated as incorrect. All suggested fix solutions (of the correct examples) were correct. The correctness ratio of the open-ended setting is 0.84, which is sufficiently high. We analyzed the background of those workers who submitted incorrect examples. Most of them use Facebook

often (40%, 40%, and 20% use Facebook every day, every week, and every month). The average self-evaluation of privacy scale of those workers who submitted incorrect examples is 2.8 (out of 5). This shows that the workers’ background knowledge indeed impact the correctness of their ideas. We also did the same measurement of the closed-ended setting. Out of thirty-six examples, thirteen examples are incorrect. The correctness ratio is 0.64. Regarding the workers who made incorrect answers, most of them use Facebook often (42%, 42%, and 16% use Facebook every day, every week, and every month). 71% of them considered themselves adequately protected their own private information on Facebook. The average self-evaluation of their privacy scale is 3.4 (out of 5). Surprisingly the workers of the closed-ended experiments are more confident in their privacy knowledge than the open-ended experiments, while they produce more incorrect answers. This shows that the general crowd may over-estimate their understanding of privacy.

Second, we measure the novelty of the ideas. We choose different similarity thresholds (we use the same threshold value of  $\delta$  and  $\theta$ ). Figure 7 shows the results for both the open-ended and closed-ended settings. All reported results are computed as the average of novelty scores of all example postings. We have the following observations. First, the novelty reduces when the similarity threshold decreases. This is not surprising as higher similarity thresholds lead to fewer similar ideas, and thus higher novelty. Second, for both settings, the novelty is sufficiently high (at least 0.88). This shows that the crowd respected the creativity requirements in the HIT description, and generated high-novelty ideas. Third, the novelty of ideas generated in the closed-setting is higher than in the open-setting. This is not surprisingly as in the open-end setting, the workers incline to generate ideas on the popular privacy topics. These ideas have high probability to be similar. But in the closed-end setting, the picked privacy topics are those that the general public rarely pay attention to (as proven by the analysis of the open-end setting results in Section 3.2). Ideas generated on those topics are less likely to collude and more likely to be dissimilar. We also asked the privacy experts to review the ideas from the crowd regarding the novelty. The privacy experts agreed that some privacy examples with high novelty scores were indeed new and urging.

Given both high correctness ratio and high novelty, we believe that in terms of game-based privacy learning, the non-expert crowd indeed can generate high-quality game design ideas from their experiences, as they may be able to generate ideas in new ways and may have access to solutions that the experts do not.




Figure 7: Novelty of Crowd-generated Ideas

	Privacy expertise	Privacy experience
Correctness	0.42	-0.05
Novelty	-0.22	-0.32

Table 1: Correlation between idea quality and Workers’ background


**Correlations.** To investigate whether there exists a relationship between the background of an individual and the quality of his generated ideas, we measure the Pearson correlation between the correctness & novelty of ideas and the workers’ privacy background & experiences. We consider the workers’ self-evaluation of expertise level as their privacy expertise, and their self-evaluation of whether having provided sufficient privacy protection on social networking sites as their privacy experience. The correlation results are shown in Table 1. We have the following findings. First, the idea correctness is positively correlated to the workers’ privacy expertise. This is not surprising as with more expertise the workers understand the problem better and thus produce more suitable answer. Second, the correlation between the idea correctness and the workers’ online privacy experience is very weak. This is somewhat surprising but still can be explained: Getting more experiences may not be able to help the workers to understand privacy better. Sometimes the users repeat the same mistakes if they were not aware. Regarding the novelty of ideas, the correlation between the novelty and the expertise, as well as the correlation between the novelty and experience, were negative. This implies that indeed the workers’ background does not help to generate novel ideas. We must admit that the correlation results may be biased, as the workers may over-estimate their expertise level. They may also be over-confident of their online privacy experience.

## 6.5 Game Evaluation


The game evaluation involves both crowd workers on Amazon Mechanic Turk (AMT) and three privacy experts. One hundred and twenty AMT workers participated into the evaluation task. Each worker is required to finish at least one game character, and answered all questions in the survey based on the character(s) that (s)he played with. The survey questions asked the workers to evaluate both the usability and the education success. We ensure that all characters were reviewed by the same number of workers by assigning four surveys (one for each character) to the same number of workers.

**Detecting cheating workers.** We record the time that each participant took to finish the survey (including game playing). It turns out that the length of the working time is dramatically diverse. Figure 8 shows the distribution of the time length. The length of the working time varies from 2.78 minutes to 33.9 minutes. The average of working time length is 8.3 minutes. As our privacy experts spent two minutes at average to finish playing one game character, we expect that the survey should take at least four minutes to finish (including game playing). We consider those workers who finished the survey in less than four minutes as *cheating*. There are 10 (out of 120) cheating workers based on the time analysis. Therefore, we consider our collected evaluation results acceptable.






(a) Usability evaluation: game interface



(b) Usability evaluation: game content



(c) Evaluation of education success

Figure 9: Game Evaluation




Figure 8: Time to Finish Survey

**Tool usability.** The tool usability evaluation is performed in the format of survey. The survey includes the questions that ask users to score on both the game interface and the content (More details of the survey questions are in Section 5). The evaluation scores are in a scale of 1 to 5, with 1 being strongly disagree, and 5 being strongly agree. We compute the average of evaluation scores per evaluation component per character. Figure 9 (a) and (b) show the results of tool usability evaluation. Most of the components received the evaluation score no lower than 3.03. In particular, the crowd are satisfied with the game interface, with the evaluation score no lower than 4, for all four characters (Figure 9 (a)). The crowd also enjoyed the examples and the suggested fix solutions (scores at least 3.36), for all four characters (Figure 9 (b)). We compute the average score of the game interface and the content. The Sophia character receives the highest average score, while the Olivia character receives the lowest score.

**Education success.** The education success evaluation includes the grading of the four learning components, namely, reaction measurement, learning measurement, behavior measurement, and results measurement. The score scale is the same as those questions for tool usability evaluation. The evaluation results are shown in Figure 9 (c). We compute the average score of the four components. It turned out that the average education score is in the range of [3.23, 3.83]. In other words, the crowd considers the game educational in general. From the scores of each individual component, we observe that the crowd agrees that the game works the best on the reaction measurement, but the worst on the behavior measurement. Regarding the education impact of different game characters, the Aiden character receives the highest average score for education, while the Olivia character receives the lowest score.

**Crowd VS. experts.** We asked the privacy experts to play the game and finish the same survey as the AMT workers. We calculated the average scores of both usability and education success by the experts, and computed the difference between the average scores between the experts and the crowd. It turns out the experts’ opinions are close to the crowd in terms of tool usability (average score difference  $< 0.3$ ). The experts also agree that the game is educational in general, but with scores 15% lower than the crowd. Surprisingly, the experts picked the character Lucas as the most educational, but Aiden the worst. We interviewed with the experts and found that the disagreement comes from some questions that the crowd labeled as inspiring but considered as trivial by the experts. This proves that indeed it is possible that the privacy experts may have different opinions from the general public regarding the effectiveness of the education game.

## 7. RELATED WORK

Using the Internet and related technologies, crowdsourcing is able to access the brainpower of many people and achieve a wide range of outcomes, from the simple and mundane task of collecting street addresses of companies, to more sophisticated achievement, such as Wikipedia, innovation competitions, and helping solve cutting-edge scientific problems [12, 17]. One type of tasks that are popular in the field of computer science is to combine human-generated results with computer algorithms to achieve synergistic outcomes [17]. For example, crowds’ input has been used to improve automatic extraction of concepts in texts [11] and to improve search results from search engines [8]. This type of crowdsourcing builds the foundation for more advanced techniques for large scale human-computer collaboration.

Another popular type of crowdsourcing is to give creative tasks to a crowd. In this type, instead of having people generate some close-ended answers, people are explicitly told to generate novel outcomes. Yu and Nickerson [35] conducted studies in which a crowd developed chair design sketches and successive crowds combined the designs of previous crowds, resembling genetic algorithm in computer science. The authors showed that crowd members tended to integrate novel and practical features of design, which help improve creativity [36]. In another study, crowd-based genetic algorithm was used to generate graphical advertisements of an online game. The results showed that having crowds modifying previous generation of ads generated better ads than having crowds combining previous ads [29].

In addition to making design sketches, textual idea generation is also a popular creative task in crowdsourcing re-

search. An important question is whether crowds are able to generate ideas of similar quality as professionals. Poetz and Schreier [27] show that product ideas from an idea contest with customers have higher novelty and customer benefit than professionals' ideas, although the idea feasibility is somewhat lower. Some top ideas that the executives like are from customers, or crowds. Therefore, it is possible for crowds to generate useful ideas, especially if they are the end users of the products. Some specific technique in crowdsourcing idea generation show their effectiveness in improving idea creativity, such as deliberately finding source of analogies from other websites [34] and decomposing the initial creative task into sub-problems [23].

In addition to idea generation, idea evaluation is also an important topic in crowdsourcing creative tasks. It is found that using multiple scales (e.g., novelty, value, feasibility) to measure idea quality is beneficial if ideas are not long [30]. If ideas are long, a single scale measurement (only idea quality) may lead to more accurate evaluation. The study also finds that having 20 ratings for an idea generates accurate evaluation. In another study of crowd-generated ideas, it is found that prediction voting (predicting whether an idea can win the competition) is more appropriate when many poor ideas need to be filtered out, while Likert scale ratings are more appropriate when more refined distinctions need to be made for ideas that are of reasonable quality [6].

## 8. CONCLUSION

In this paper, we present our effort of leveraging crowdsourcing techniques for game-based privacy learning. We utilize the crowd for both idea collection for the game design and game evaluation. Our experiments show that the non-expert crowd can provide useful inputs for both game design and evaluations.

For the future work, we plan to improve *DEPET* based on the received feedback from the *game evaluation* phase. We also plan to explore if under certain assumptions the "wisdom of the crowd" is able to outperform a smaller group of "experts". Longer term, we want to investigate the potential of recruiting non-expert workers to collaborate on designing large-scale content-creation projects.

## 9. ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant # SaTC-1464800. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## 10. REFERENCES

- [1] Amazon mechanic turk. <https://www.mturk.com/mturk/welcome>.
- [2] The sheep market. <http://www.thesheepmarket.com/>.
- [3] Social networking fact sheet. <http://www.pewinternet.org/fact-sheets/social-networking-fact-sheet/>.
- [4] A. Albrechtslund. Online social networking as participatory surveillance. *First Monday*, 13(3), 2008.
- [5] O. Alonso, D. E. Rose, and B. Stewart. Crowdsourcing for relevance evaluation. In *ACM SigIR Forum*, volume 42, pages 9–15, 2008.
- [6] J. Bao, Y. Sakamoto, and J. V. Nickerson. Evaluating design solutions using crowds. In *AMCIS 2011 Proceedings - All Submissions. Paper 446*. Association for Information Systems, 2011.
- [7] D. M. Blei, A. Y. Ng, and M. I. Jordan. Latent dirichlet allocation. *the Journal of machine Learning research*, 3:993–1022, 2003.
- [8] A. Bozzon, M. Brambilla, and S. Ceri. Answering search queries with crowdsearcher. In *Proceedings of the 21st international conference on World Wide Web*, pages 1009–1018, 2012.
- [9] D. C. Brabham. Moving the crowd at threadless: Motivations for participation in a crowdsourcing application. 2009.
- [10] F. P. Brooks Jr. *The design of design: Essays from a computer scientist*. Pearson Education, 2010.
- [11] G. Demartini, D. E. Difallah, and P. Cudré-Mauroux. Zencrowd: leveraging probabilistic reasoning and crowdsourcing techniques for large-scale entity linking. In *Proceedings of the 21st international conference on World Wide Web*, pages 469–478, 2012.
- [12] C. B. Eiben, J. B. Siegel, J. B. Bale, S. Cooper, F. Khatib, B. W. Shen, F. Players, B. L. Stoddard, Z. Popovic, and D. Baker. Increased diels-alderase activity through backbone remodeling guided by foldit players. *Nature biotechnology*, 30(2):190–192, 2012.
- [13] E. Hargittai and D. Boyd. Facebook privacy settings: Who cares? *First Monday*, 15(8), 2010.
- [14] Y. Huang, Y. Wang, and Y. Tang. Privacy in emotion sharing on social media. In *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, 2014.
- [15] G. Kazai, J. Kamps, M. Koolen, and N. Milic-Frayling. Crowdsourcing for book search evaluation: impact of hit design on comparative system ranking. In *Proceedings of the 34th international ACM SIGIR conference on Research and development in Information Retrieval*, pages 205–214, 2011.
- [16] D. L. Kirkpatrick. Techniques for evaluating training programs. *Classic writings on instructional technology*, 1:231–241, 1979.
- [17] A. Kittur, J. V. Nickerson, M. Bernstein, E. Gerber, A. Shaw, J. Zimmerman, M. Lease, and J. Horton. The future of crowd work. In *Proceedings of the 2013 conference on Computer supported cooperative work*, pages 1301–1318, 2013.
- [18] A. Kittur, B. Smus, S. Khamkar, and R. E. Kraut. Crowdforge: Crowdsourcing complex work. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, pages 43–52, 2011.
- [19] N. Koudas, S. Sarawagi, and D. Srivastava. Record linkage: similarity measures and algorithms. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 802–803, 2006.
- [20] S. H. Lee. Usability testing for developing effective interactive multimedia software: Concepts, dimensions, and procedures. *Educational Technology & Society*, (2), 1999.
- [21] J. M. Leimeister, M. Huber, U. Bretschneider, and H. Krcmar. Leveraging crowdsourcing:

- activation-supporting components for it-based ideas competition. *Journal of management information systems*, 26(1):197–224, 2009.
- [22] H. Li and Y. Sakamoto. The influence of collective opinion on true-false judgment and information-sharing decision. In *Annual Meeting of the Cognitive Science Society*. Cognitive Science Society, 2013.
- [23] L. Luo and O. Toubia. Improving online idea generation platforms and customizing the task structure based on consumers’ domain specific knowledge. *Journal of Marketing*, 2015.
- [24] H. McCracken. Teaching with technology: Tools and strategies to improve student learning. *Faculty Focus*, 2011.
- [25] C. H. Papadimitriou, H. Tamaki, P. Raghavan, and S. Vempala. Latent semantic indexing: A probabilistic analysis. In *Proceedings of the seventeenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 159–168, 1998.
- [26] F. Pass and O. Firssova. *Usability Evaluation of Integrated E-Learning*. New York: RoutledgeFalmer (Taylor & Francis Group), 2004.
- [27] M. K. Poetz and M. Schreier. The value of crowdsourcing: can users really compete with professionals in generating new product ideas? *Journal of Product Innovation Management*, 29(2):245–256, 2012.
- [28] M. Prensky. Computer games and learning: Digital game-based learning. *Handbook of computer game studies*, 18:97–122, 2005.
- [29] J. Ren, J. V. Nickerson, W. Mason, Y. Sakamoto, and B. Graber. Increasing the crowd’s capacity to create: how alternative generation affects the diversity, relevance and effectiveness of generated ads. *Decision Support Systems*, 65:28–39, 2014.
- [30] C. Riedl, I. Blohm, J. M. Leimeister, and H. Krcmar. The effect of rating scales on decision quality and user attitudes in online innovation communities. *International Journal of Electronic Commerce*, 17(3):7–36, 2013.
- [31] Y. Sakamoto and J. Bao. Testing tournament selection in creative problem solving using crowds. In *Proceedings of the International Conference on Information Systems (ICIS)*, 2011.
- [32] N. Shiratuddin and M. Landoni. Evaluation of content activities in children’s educational software. *Evaluation and Program Planning*, 25(2):175–182, 2002.
- [33] A. Sisarica and N. Maiden. An emerging model of creative game-based learning. In *Proceedings of the International Conference of Serious Games Development and Applications (SGDA)*, pages 254–259. 2013.
- [34] L. Yu, A. Kittur, and R. E. Kraut. Searching for analogical ideas with crowds. In *Proceedings of the 32nd annual ACM conference on Human factors in computing systems*, pages 1225–1234, 2014.
- [35] L. Yu and J. V. Nickerson. Cooks or cobblers?: crowd creativity through combination. In *Proceedings of the SIGCHI conference on human factors in computing systems*, pages 1393–1402, 2011.
- [36] L. Yu and Y. Sakamoto. Feature selection in crowd creativity. In *Foundations of Augmented Cognition. Directing the Future of Adaptive Systems*, pages 383–392. 2011.
- [37] N. Zeichner, J. Berant, and I. Dagan. Crowdsourcing inference-rule evaluation. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics: Short Papers-Volume 2*, pages 156–160, 2012.

This figure "HITInstruction1.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>



This figure "edu-success.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "evalGUI.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "example.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "frequency.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>



This figure "gameContent-eval.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "gameGUI-eval.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "groupfreq.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "novelty-close.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>



This figure "novelty.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "posting.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "postingIns.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "privacy-level.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>



This figure "self-privacy.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "set2-frequency.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "set2-plevel.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "set2-scale.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "solution.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "solutionIns.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "topicfreq.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>

This figure "workTime.png" is available in "png" format from:

<http://arxiv.org/ps/1603.02766v1>